Generating Interest Rate Scenarios A Real-World Approach

By Stephen J. Strommen FSA, MAAA, CERA

Copyright 2023

Permission is hereby granted to copy and distribute this document freely but not to change it in any way.



Contents

Preface
Chapter 1: Basic theory of short-rate models of the yield curve4
Chapter 2: A general class of short rate models6
Chapter 3: Stylized facts and their role in model design7
Design step 1: A simple random walk7
Design step 2: A one-factor mean-reverting model8
Design step 3: Skewing the distribution of future interest rates
Design step 4: Stochastic volatility10
Design step 5: Hump and valley yield curves – the mean reversion process11
Design step 6: Pinning the short rate at zero13
Chapter 4: Fine-tuning the model14
Stretching the tails14
Eliminating extremely negative slopes when interest rates are high16
Chapter 5: The scenario generation process17
Completing the yield curve
Generating the database19
Specifying the contents19
Calculating the yield curves19
Retrieving a yield curve21
Initializing the state variables22
Generating scenarios using the database23
Appendix 1: Summary of stochastic processes in this model
Volatility process
Mean reversion point process
Short rate process
Updating the state variables for each time step27
Appendix 2: Suggested calibration and illustrative scenario paths
Adjusting for the discrete time step
Volatility process
Mean reversion point process
Short rate process

Preface

Stochastic economic scenario generators are widely used in the financial services industry. Among investors they are used for valuation of complex derivative securities often used for hedging risk. In the insurance industry they are used both for risk management and for financial reporting that requires valuation of contracts involving risk. Stochastic scenario generators are used to answer questions like "How bad could the financial outcome be", and "How likely is that?", and to put a value on financial contracts with an uncertain and risky outcome.

The work presented in this book was motivated by a development in insurance industry regulation in the U.S.A. The National Association of Insurance Commissioners (NAIC) sets minimum standards for the total amount of assets insurers must hold for payment of future insurance claims, an amount which by its nature is uncertain. In some cases, those requirements are based on analysis using stochastic economic scenarios. The NAIC mandates the stochastic economic scenario generator that must be used in the analysis that sets minimum total asset requirements.

In years following the 2008 financial crisis, the NAIC recognized that the economic scenario generator in use could not generate scenarios containing economic conditions as extreme as those experienced in the financial crisis. The generator needed update or revision to reflect the likelihood that such a crisis could happen again in the future.

Updating the generator for this purpose proved to be both difficult and controversial. Stochastic models that reproduce some of the behaviors seen in the crisis did not exist. Crude adaptations of existing models were proposed, but many observers found them unsatisfactory. A national field test of one proposed new generator failed with unsatisfactory results.

The work presented in this book was undertaken by the author to create a new stochastic economic scenario generator that could be used by the NAIC for this purpose. Due to considerations outside the scope of this book, the generator presented here was never seriously considered by the NAIC. Nevertheless, it could be useful to insurance risk managers and others. Some new scenario generation methods, not previously widely used, are employed.

The focus of this book is on the interest rate generator, which produces scenarios of yield curves for US Treasury bonds. This proved to be the most controversial part of the generator used by the NAIC¹. The generator in use did not produce scenarios where the short-term interest rate was pinned at zero, as happened during the financial crisis, and it was felt that the generator should produce some such scenarios, possibly including negative interest rates. Even at rates above zero there was concern among

¹ The generator used by the NAIC is not limited to yield curves for US Treasury bonds. It includes credit spreads and default rates for corporate bonds and includes returns for various classes of equity investments. The generation of these additional variables is not addressed in this book.

regulators that the generator did not produce enough "low-for-long" scenarios. In addition, the generator in use was neither arbitrage-free nor consistent with market pricing of fixed income securities, which are prevalent in life insurance company investment portfolios. The author felt that any new generator should have these properties, since the underlying theory is now widely understood and accepted.

This book is written to introduce a new interest rate scenario generator and to explain a new scenario generation process. The new generator is a short-rate model, where the yield curve is based on evolving expectations concerning the future path of the short-term rate. Much research has been focused on short-rate models for which one can find a closed-form mathematical expression to complete the yield curve. The model presented here does not have such a closed-form expression. A new scenario generation process is presented to deal with the practical problem this presents.

A central problem in this kind of work is that the historical data is a limited sample. The period from the Great Depression (1930's) until today covers a small number of complete economic cycles. A scenario generator is used to create thousands of such cycles. It is not possible to find statistically significant evidence to support every feature of a generator that may affect the fit with historical data because the data set is too small. Therefore, significant reliance is placed on judgment and belief in an underlying theory of market behavior. Disagreements concerning theory and judgment can arise even among experts in the field. Several potential disagreements with the theory and judgments employed in this work are highlighted in the book for educational purposes.

The first three chapters outline the theory underlying short-rate models, the mathematical form of the class of models being considered, and the stylized facts that guide model design and calibration. Chapter 4 presents some original and unique approaches to fine-tuning the model to better reflect some aspects of historical behavior. Chapter 5 presents a practical scenario generation process for use when completion of the yield curve requires numerical simulation.

Stephen J. Strommen Fall 2023

Chapter 1: Basic theory of short-rate models of the yield curve

The "short rate" is the rate of interest for a very short period such as overnight. The yield curve is a vector of interest rates for funds locked in for longer periods of time. A yield curve has a rate for each length of time T for which funds may be locked in. The "risk-free" yield curve defines the price today of an amount of cash to be received at time T in the future, with no risk of default.

Defining terms:

 r_t = the short rate at time t

 $r_{t,T}$ = the rate at time t for a payment to be received at time T.

The yield curve at time t is a set of values $r_{t,T}$ for various values of T. The notation for the short rate r_t is just shorthand for $r_{t,0}$.

The price at time t of a payment at time t + T is $P_{t,T} = (1 + r_{t,T})^{-T}$. With a set of observed market prices at time t, one can calculate the market yield curve using the expression $r_{t,T} = P_{t,T}^{(-1/T)} - 1$.

A short-rate model is one where the yield curve is connected to the expected future path of the short-term rate.

In a market-consistent risk-neutral model (where the market price of risk is zero), the yield curve is defined as follows:

$$P_{t,T} = E^{Q} \left[exp\left(-\int_{t}^{t+T} r_{s} ds \right) \right]$$

Note that the expectation is taken using the Q measure, or risk-neutral probabilities, which are not real probabilities.

A model where the expectation is taken using the P measure, or real-world probabilities, is called a market-coherent model. In a market-coherent real-world model, the market price of risk is not zero. Prices of securities that involve risk are lower than prices of securities with no risk. The prices $P_{t,T}$ for T > t involve risk due to the volatility of the short—term rate and the resulting volatility of the entire yield curve. The market price of risk reduces $P_{t,T}$ and increases $r_{t,T}$ for values of T > t. The increase in $r_{t,T}$ is called a term premium and we will denote it as $\varpi_{t,T}$. Note that values of $\varpi_{t,T}$ increase with increasing T.

The market price of risk is commonly defined as follows:

$$\lambda = \frac{(Yield \ premium \ over \ the \ risk - free \ rate)}{(Standard \ deviation \ of \ price)}$$

In this model λ is a parameter with a constant value roughly in the range 0.2 to 0.3². The term premium is the expected yield premium and is equal to λ times (Standard deviation of price). Since the price $P_{t,T}$ is a function of the short rate, the standard deviation of the price depends on the volatility of the short rate and can be calculated or estimated numerically.

(Standard deviation of price) = (Standard deviation of
$$r_t$$
) $\frac{\partial P_{t,T}}{\partial r_t}$

That means, in turn, that values of the term premiums $\varpi_{t,T}$ can be calculated.

$$\varpi_{t,T} = \lambda$$
(Standard deviation of price)
 $\varpi_{t,T} = \lambda$ (Standard deviation of r_t) $\frac{\partial P_{t,T}}{\partial r_t}$

With those term premiums, we can define the price in a market-coherent real-world model as follows:

$$P_{t,T} = E^{P}\left[exp\left(-\int_{t}^{t+T} (r_{s} + \varpi_{t,T})ds\right)\right]$$

Note that the expectation is taken using the P measure, or real-world probabilities. That makes this a real-world model.

The challenge in building a short rate model of the yield curve is to define a stochastic process for the short rate that leads to realistic behavior of not only the short rate but also the entire yield curve when the yield curve $r_{t,T} = P_{t,T}^{(-1/T)} - 1$ is defined by the spot prices $P_{t,T}$ based on the formula above.

² This is the first of many issues on which opinions differ among experts in this field. Some believe the market price of risk is not constant over time, and some believe that it is not constant for all values of T at any point in time t. Since the market price of risk is not directly observable, this question cannot be resolved conclusively. Cross-sectional studies over time have come to different conclusions, sometimes due to differences in the time step over which observations of price volatility or yield premiums are measured. The reasoning used for this model is that the market price of risk is certainly nonzero in the real world, and a constant value does not unnecessarily complicate the model while it does lead to reasonable term premiums. We use a time step of 1 month in this work because that is the time step used by the current model mandated by the NAIC for this purpose.

Chapter 2: A general class of short rate models

For our purposes, we express a general short-rate model in discrete form as follows:

$$r_t = r_{t-1} + d(r_{t-1}) + s(r_{t-1})$$

Where:

d(r) = a drift function; the movement due to expected drift in one time step

s(r) = a random shock function; the random movement due to volatility in one time step

A short-rate model has parameters and one or more state variables called "factors". The definition above is for a one-factor model, and the state variable for that one factor is the short-term rate itself. The state of that factor at any time step is the value of the short-term rate. The parameters are used inside the functions d(r) and s(r). For example, the Vasicek model in discrete form defines d(r) and s(r) as follows:

 $d(r) = \alpha(\theta - r)$ θ is the mean reversion point parameter α is the mean reversion speed parameter

 $s(r) = \sigma X$ σ is the volatility parameter X is a random number from a normal distribution with mean zero and variance one

Either of these functions may depend on other factors besides the short rate r. Introduction of other factors creates a multi-factor model. For our purpose here, we introduce another factor into each of d(r) and s(r) to make this a three-factor model, that is, a model with three state variables.

The drift function becomes $d(r_t, c_t)$ where c_t is a moving central mean reversion point.

The random shock function becomes $s(r_t, v_t)$ where v_t is a moving volatility.

The result is a generalized double-mean-reverting model with stochastic volatility.

$$r_t = r_{t-1} + d(r_{t-1}, c_{t-1}) + s(r_{t-1}, v_{t-1})$$

To completely specify the model, we need to define the stochastic process for each of the additional state variables c, v and define the functions d(r, c) and s(r, v) used in the stochastic process for the state variable r. These definitions are introduced later in this paper and summarized in Appendix 1.

A model like this is fitted to a starting yield curve by selecting starting values for the state variables r_t , c_t , v_t so that the yield curve based on them closely matches the yield curve observed on the start

date (time *t*). No model ever matches exactly, but a three-factor model can closely mimic a variety of starting yield curve shapes.

Chapter 3: Stylized facts and their role in model design

As mentioned earlier, the challenge in building a short rate model of the yield curve is to define a stochastic process for the short rate that leads to realistic behavior of not only the short rate but also the entire yield curve. With that in mind, it is useful to stipulate what realistic behavior looks like. The following stylized facts represent aspects of observed behavior of the yield curve.

- 1. Interest rates fluctuate in a range of roughly 0% to 20% or so.
- 2. The short-term end of the yield curve tends to be more volatile than the long-term end.
- 3. The yield curve slope is more often upward when interest rates are low and downward when interest rates are very high.
- 4. The probability distribution of future interest rates is skewed with a longer tail on the high side. The median is below the mean.
- 5. Volatility tends to be higher when interest rates are higher and vice versa.
- 6. The degree of volatility varies over time, and not just by the level of interest rates.
- 7. The yield curve can sometimes have a hump or valley shape.
- 8. The Federal Reserve has so far imposed a hard floor of zero on the short-term rate. When the short-term rate is floored at zero, the long end of the yield curve continues to fluctuate.

In designing our model, we start with a simple random walk for the short rate and add complexity only as needed to simulate the stylized facts listed above.

Design step 1: A simple random walk

The first stylized fact says interest rates fluctuate. We can make them fluctuate using a random walk where random shocks in each time step move the rate up or down.

A simple random walk fits within the class of models we've defined. The stochastic process is:

$$r_t = r_{t-1} + d(r_{t-1}, c_{t-1}) + s(r_{t-1}, v_{t-1})$$

Where:

$$d(r_{t-1}, c_{t-1}) = 0$$

 $s(r_{t-1}, v_{t-1}) = \sigma X_t$

 σ is a constant volatility parameter

 X_t is a random number for time step t drawn from a normal distribution with mean zero and variance one.

The range of future values produced by a random walk like this is unlimited. The standard deviation of future values *T* time steps in the future is $\sigma\sqrt{T}$. This is problematic because the first stylized fact says that interest rates stay in a limited range.

Design step 2: A one-factor mean-reverting model

To constrain the range of values produced by our model, we introduce a tendency for the value to drift towards a central mean.

Instead of $d(r_{t-1}, c_{t-1}) = 0$ as in step 1, define the drift function as follows:

 $d(r_{t-1}, c_{t-1}) = \theta(\mu - r_{t-1})$ μ is a constant mean reversion point parameter θ is a constant mean reversion speed parameter

The drift function creates a downward drift when the interest rate is above the mean and an upward drift when the interest rate is below the mean. The drift gets stronger as the interest rate gets farther from the mean. The limits the range of interest rates generated, and the range depends on the parameters μ and θ as well as the volatility σ .

Introducing mean reversion has an affect not only on the short rate but also on the behavior of the entire yield curve in a short rate model. It can be shown that behavior of the yield curve is consistent with stylized facts 2 and 3:

- 2. The short-term end of the yield curve tends to be more volatile than the long-term end.
- 3. The yield curve slope is more often upward when interest rates are low and downward when interest rates are very high.

This one-factor mean-reverting model is a discrete version of the Vasicek model, which is sometimes used due to its simplicity and its ability to simulate the first 3 stylized facts of yield curve behavior.

Design step 3: Skewing the distribution of future interest rates

Stylized facts 4 and 5 are:

- 4. The probability distribution of future interest rates is skewed with a longer tail on the high side.
- 5. Volatility tends to be higher when interest rates are higher and vice versa.

In the model from step 2, volatility is a constant parameter. It can be shown that if you make volatility an increasing function of the short rate, the probability distribution of future interest rates will become skewed with a longer tail on the high side and a median below the mean.

In step 1 we defined the random shock function as follows:

$$s(r_{t-1}, v_{t-1}) = \sigma X_t$$

 σ is a constant volatility parameter

 X_t is a random number from a normal distribution with mean zero and variance one, drawn for time step t.

We can enhance this by making the volatility parameter a volatility function:

v(r) = volatility as a function of the short rate.

We then define the random shock function as $s(r_{t-1}, v_{t-1}) = v(r_{t-1})X_t$

This is the approach embedded in some widely used short rate models. The Cox-Ingersoll-Ross (CIR) model uses $v(r) = \sigma \sqrt{r}$ where σ is a constant parameter. The volatility approaches zero as the short rate approaches zero. In the continuous form of the CIR model this prevents the short rate from ever getting down to zero because mean reversion takes over when volatility vanishes³.



Historical data are available to visualize how volatility depends on the level of interest rates.

The blue dots represent historical observations, and the orange line represents a reasonable central value at each level of the interest rate. For our model we define v(r) to be represented by the orange line or a similar line that reasonably represents the data. Values on that line are tabulated and interpolated as necessary for any value of r.

We take this approach for these reasons:

³ In the discrete form of the CIR model, it is possible to generate a short rate below zero, but this is typically corrected by imposing a zero floor. If not corrected, the volatility for the next time step is undefined because $\sigma\sqrt{r}$ involves the square root of a negative number.

- It is more closely tied to historical observations. The shape of the orange line does not resemble either a square root curve or other mathematical function as used in other recognized models. The reason those models used such formulations was a requirement that there be a formulaic way to calculate the entire yield curve. The models had to be "mathematically tractable". In our model we dispense with that requirement and always calculate the yield curve through numerical simulation.
- We observe that volatility does not approach zero as the short rate approaches zero. In our limited history, it has abruptly dropped to zero when the short rate is pinned at zero, but not when the rate is above zero by even a little. We will address behavior when the short rate is pinned later in this book.

Design step 4: Stochastic volatility

Stylized fact number 6 says:

6. The degree of volatility varies over time, and not just by the level of interest rates.

Volatility tends to increase during periods of economic uncertainty, and that can occur regardless of the level of interest rates. Volatility represents a risk that financial institutions must manage, so our model would not fully serve its purpose if it did not simulate periods of above-average volatility.

We will introduce this behavior with a volatility multiplier whose value is centered on 1.0 and which fluctuates above and below 1.0.

There are many ways this could be done. Here we make the stochastic multiplier follow its own meanreverting random walk, with a mean reversion point of 1.0. Stochastic volatility tends to have a skewed probability distribution; it cannot go below zero and maximum volatility is more than double normal volatility, so volatility has a longer tail on the high side than the low side. To reflect this, we model the multiplier as a mean-reverting random walk with a symmetric distribution (like the Vasicek model) and then transform the result using the exponential function.

Let x_t be the random walk value of the multiplier at time t.

$$x_{t+1} = x_t + \gamma(\mu - x_t) + \sigma X_t$$

Parameters:

 $\mu =$ Mean reversion point for the volatility multiplier (normally 1.0 or near 1.0)

- $\sigma = -$ Volatility of the multiplier
- $\gamma =$ Mean reversion strength for the multiplier

Random element at each time step:

 X_t = Random shock (standard normal distribution, mean 0, variance 1)

Since the distribution of volatility tends to be skewed with a long tail on the high side, use an exponential transformation of this multiplier $(e^{x_{t-1}-1})$ to calculate volatility for time step t. If we define v(r) to be the tabulated volatility for interest rate r we have:

Stochastic volatility = $v_t = v(r)e^{(x_{t-1}-1)}$

This transformation $(e^{x_{t-1}-1})$ retains the central value of 1 when $x_{t-1} = \mu = 1$ but makes increases larger and decreases smaller. It also prevents the multiple from ever going below zero. When calibrating the stochastic volatility process, one needs to be aware that the mean of the transformed multiplier $e^{x_{t-1}-1}$ is larger than the mean reversion point parameter μ used in the stochastic process. The transformed value has a slightly larger mean.

The random shock term in our model is the following.

$s(r_{t-1}, v_{t-1}) =$	$v_{t-1}X_t = v(r_{t-1})e^{x_{t-1}-1}X_t$ where:
$v(r_{t-1})$	is the tabulated volatility for the short rate at the end of the prior time step
X_{t-1} X_t	is a random shock (standard normal distribution, mean 0, variance 1).

The result is a stochastic log volatility model, even though it has not been expressed in the typical way. To see this, note that the coefficient of X is the volatility. If we use v to denote the volatility, we have the following:

$$v = v(r)e^{x-1} = e^{\log(v(r)) + x-1}$$

In the exponent on the right-hand side, the stochastic process for x - 1 (which is centered on zero) is added to the log of the tabulated volatility. The resulting exponent is a stochastic process centered on the log of the tabulated volatility. That makes this a stochastic log volatility process; we exponentiate a stochastic process for the log of the volatility.

Design step 5: Hump and valley yield curves – the mean reversion process Stylized fact 7 says:

7. The yield curve can sometimes have a hump or valley shape.

The model, as developed so far, cannot produce yield curves with a hump or valley shape. The yield curve slope⁴ depends only on the current value of the short rate. The lower the short rate, the greater

the slope, and vice versa. This behavior results from having fixed parameter values for the mean reversion point and mean reversion strength.

If we make the mean reversion point itself a mean-reverting stochastic process that moves up and down around an ultimate long-run mean, the model can produce yield curves with a hump or valley shape.

- A humped curve occurs when the short rate is less than the current mean reversion point (so it initially drifts upward) and the current mean reversion point is above the ultimate mean reversion point (so it drifts downward over time).
- A valley-shaped curve occurs when the short rate is greater than the current mean reversion point (so it initially drifts downward) and the current mean reversion point is below the ultimate mean reversion point (so it drifts upward over time).

There are many ways to add a moving mean reversion point to the model. The chosen method is purely a matter of judgment. The purpose here is to introduce this feature into the model, not to suggest an optimal way to do it. The choice of approach is based on a generous helping of intuition and judgment.

For this model the current mean reversion point is a state variable with a stochastic process like the one-factor Vasicek model for interest rates, but with these modifications:

- **Correlation of shocks.** The random shocks to the current mean reversion point are correlated with random shocks to the short-term rate. The correlation rate is an additional parameter.
- Volatility is tabulated. In the Vasicek model volatility is constant. In this model the volatility is tabulated by the level of the current mean. In the suggested calibration, volatility is smaller at lower interest rates but remains positive if the current mean goes below zero to some degree. See Appendix 2 for the suggested calibration.
- Hard floor. There is a hard floor to the value of the current mean at the end of any time step. If the stochastic process would produce a value less than the floor, it is set equal to the floor. The combination of nonzero volatility and a floor adds a sort of bounce-back behavior resulting in many values just above the floor.

For this model, the moving mean reversion point c_t is defined in the following way.

$$c_{t} = \max \{ f, c_{t-1} + \gamma(\theta - c_{t-1}) + \sigma(c_{t-1})X_{t} \}$$

 θ = a fixed long-term mean around which the moving mean fluctuates.

 γ = the rate at which the moving mean is attracted to (reverts to) the fixed long-term mean.

 $\sigma(c_{t-1}) =$ volatility tabulated by the value of c and interpolated as necessary.

$$f = floor$$

 X_t = Random shock (standard normal distribution, mean 0, variance 1)

In a model like this where the moving mean reversion point is a state variable, the initial value of the moving mean reversion point c_t is not set based on actual history. State variables are initialized by choosing values that best fit the shape of the starting yield curve defined by spot prices P_T .

The moving mean reversion point is used as a state variable in our short rate model. The drift term d(r, c) in our short rate model uses the moving mean reversion point discussed above as its second parameter. d(r, c) is defined as follows:

$$d(r,c) = \alpha(c-r)$$

 α = the rate at which the short rate is attracted to (reverts to) the current mean reversion point.

Design step 6: Pinning the short rate at zero

Stylized fact number 8 says:

8. The Federal Reserve has so far imposed a hard floor of zero on the short-term rate. When the short-term rate is floored at zero, the long end of the yield curve continues to fluctuate.

The Federal Reserve (the Fed) largely controls short-term interest rates in the United States. They use this power to govern the economy with a particular focus on unemployment and inflation. When the economy overheats and inflation becomes a problem, the Fed raises interest rates to make borrowing more expensive. When the economy cools off too much and unemployment rises, the Fed reduces interest rates to stimulate investment and borrowing to increase economic activity and employment.

When unemployment rises due to economic stress, the Fed has determined there is a limit on how far they can reduce the short-term interest rate. Zero is the lower limit they have applied. At times, they have left the short-term rate pinned at zero for an extended period. During that time, the remainder of the yield curve, which they do not directly control, continues to fluctuate.

To simulate this behavior in a short rate model, we adopt what has been called the "shadow rate" approach. The shadow rate is the short-term rate that the Fed would set in the absence of the zero floor. The actual short rate is the maximum of the shadow rate and zero. When completing the yield curve using the pricing formula below, the variable r_t is the floored actual short rate, not the shadow rate.

$$P_{t,T} = E^{P}\left[exp\left(-\int_{t}^{t+T} (\max\left(0, r_{s}\right) + \varpi_{t,T}\right)ds\right)\right]$$

As mentioned earlier, much research has focused on models for which the expression above is "mathematically tractable", meaning there is a closed form formula for calculating the expectation. When the short rate r has a floor, the expression becomes intractable and must be evaluated numerically. That has been viewed as a practical barrier to using the shadow rate approach. Later in this book we show how to overcome the practical barrier. Knowing that the practical problem can be overcome, we adopt the shadow-rate approach.

Our stochastic process is then left to simulate the shadow rate, that is, the short-term rate that the Fed would set in the absence of a zero floor. In times of great financial stress, this rate could be below zero to provide the desired amount of economic stimulus. And it may stay below zero so long as that stress continues. During that time, the observed short rate would be zero and the remainder of the yield curve would continue to be governed by the formula above. Under those conditions, the mean reversion in our model will generally lead to an upward sloping yield curve, with long-term rates well above zero. Those long-term rates will continue to fluctuate with fluctuations in the shadow rate, even while it is below zero and the observed short rate is pinned at zero.

That behavior, with the short rate pinned at zero and long-term rates continuing to fluctuate, is consistent with the stylized facts of observed yield curve behavior. It is also consistent with what we know about the way the Fed manages short-term interest rates. Therefore, to conform our model to observed behavior, we adopt the shadow rate approach and introduce the zero floor to the formula above which is used to complete the yield curve.

Chapter 4: Fine-tuning the model

The model described so far generates scenarios that are consistent with the stylized facts of observed yield curve behavior. Yet there are some aspects of scenario behavior that don't quite match historical observations. In particular,

- 1. The historical record seems to have fatter tails, that is, relatively more extremely high and extremely low interest rates than this model produces.
- 2. When interest rates are very high, the model produces yield curves with more negative slopes than have been observed historically.

We can tweak the model to improve its fit to the historical record. In doing so we use intuition (sometimes called professional judgment) to suggest changes to the model that are consistent with a view of behavior by the Fed in managing the short-term rate. The justification for the tweaks introduced in this chapter may be dismissed as "hand-waving" by some, and they cannot be proven "correct". Nevertheless, they improve the fit of the model to historical data.

Stretching the tails

Some have observed that the Fed acts more strongly in responding to stress than when recovering from stress. In our model, stress is indicated by the distance between the current mean reversion point and the ultimate mean. When the current mean reversion point is further from the ultimate mean than the short rate or is on the other side of the ultimate mean, one may imply that the Fed is responding to stress rather than recovering from it. Under such conditions we can increase the rate of reversion towards the current mean.

We make the mean reversion strength larger by multiplying by a value greater than 1.0. Call the multiplier a stretch parameter. Its value can depend on the relation between the short rate, the current mean, and the ultimate mean.

In chapter 3 we defined the drift function this way:

$$d(r_{t-1}, c_{t-1}) = \alpha(c_{t-1} - r_{t-1})$$

Here we introduce a conditional multiplier η :

$$d(r_{t-1}, c_{t-1}) = \alpha \eta (c_{t-1} - r_{t-1})$$

To specify a value for η we define two parameters:

 $Z_h =$ high stretch, value ≥ 1.0

 $Z_l =$ low stretch, value ≥ 1.0

The multiplier η takes one of the following values:

If $c < \theta$ then

If
$$r > c$$
 then $\eta = Z_l$ else $\eta = 1/Z_l$

If $c > \theta$ then

If
$$r < c$$
 then $\eta = Z_h$ else $\eta = 1/Z_h$

The multiplier takes a value \geq 1.0 when either:

- $c < \theta$ and r > c (current mean below ultimate mean and short-term rate is pulled down to it)
- $c > \theta$ and r < c (current mean above ultimate mean and short-term rate is pulled up to it)

Both these conditions represent a stress to which the Fed may respond more quickly. The multiplier makes the pull to the current mean stronger in stressful conditions and weaker during recovery.

A possible justification for this is that it simulates the behavior of the Federal Reserve in adjusting shortterm interest rates. The current mean reversion point indicates whether there is a need to tighten monetary policy ($c > \theta$) or loosen it ($c < \theta$). In either case, the Fed tends to take faster action when the current mean reversion point is pulling the short rate away from the ultimate mean rather than towards it.

Experimentation has shown that using a conditional value as high as $\eta = 3$ can lead to behavior of the generated yield curves that better resembles historical behavior.

Eliminating extremely negative slopes when interest rates are high

A perennial problem with short rate models of the yield curve is that when the short rate is very high, the yield curve tends to have a steep downward slope. This occurs because mean reversion of the short rate over time becomes especially strong when the short rate is very high.

The drift term in our stochastic model is simple:

$$d(r_{t-1}, c_{t-1}) = \theta(c_{t-1} - r_{t-1})$$

This creates a strong downward drift when the short rate is very high and/or the current mean is very low. That is what causes the steep downward slope in the yield curve. One way to limit the downward drift is to prevent the current mean reversion point from getting too far below the short rate.

In chapter 3 the stochastic process for the current mean reversion point was defined as:

$$c_{t} = \max(f, c_{t-1} + \gamma(\theta - c_{t-1}) + \sigma(c_{t-1})X_{t})$$

The floor f was a constant value. To prevent c_t from getting too far below the short rate we can make the floor a function of the short rate. Experimentation has shown that an increasing function of the short rate has the desired effect on the behavior of yield curve slopes when interest rates are high.

$$f = \max(\psi, 2r_{t-1} - 0.34)$$

Here we introduce ψ to represent the floor when interest rates are low. The value of ψ would typically be zero or below. With this formula, the floor stays around ψ until the short rate rises to over 17%. Then the floor rises rapidly with the short rate. If the short rate were to reach 25%⁵, the floor on the current mean would be around 16%. The short rate and floor would tend to drift down from there, but not so fast as without the floor on the current mean reversion point.

⁵ This formula for the floor was calibrated in the context of a model with a 25% cap on the short rate. It may need revision in other models, particularly if there is significant probability of short rates over 25%. Since rates that high have never been observed in the US, such calibrations are speculative by nature.

Chapter 5: The scenario generation process

The interest rate generator is a "short rate model" that defines the stochastic process for the short-term rate. To use a short rate model in an interest rate generator, one needs a method for completing the yield curve because a short rate model on its own only provides scenarios for the short-term rate. In an arbitrage-free short rate model, the yield curve is defined by the following expression for spot prices by term to maturity T:

$$P_T = E^Q \left[exp\left(-\int_0^T r_t dt \right) \right]$$

Some short rate models are structured so that there is a formula for calculating the expression above. This requires the short rate model to be structured in a way that allows derivation of such a formula and requires, for example, that:

- any dependence of volatility on the level of the interest rate be defined by formula. For example, in the Cox-Ingersoll-Ross model the volatility is proportional to the square root of the interest rate
- all stochastic movements are based on random samples from probability distributions with no caps, floors, or truncation.

The short rate model defined here violates both of those requirements. The dependence of volatility on the level of the interest rate is tabulated rather than formulaic. And a floor is applied to paths of the short-term rate that are used in the expression above to calculate spot prices. Therefore, the yield curve must be completed using Monte Carlo simulation.

It is impractical to undertake a separate stochastic simulation to complete every yield curve in a set of stochastic scenarios. That would be a stochastic-on-stochastic process requiring too much time to be practical. An alternative process is to develop a universe (database) of available yield curves indexed by values of the state variables and select from that database (or interpolate) to retrieve the yield curve based on the state variables at each time step. Once such a database is in place, scenarios can be generated quickly. The drawback is that creation of the database can take some time. But the database only needs to be created when parameter values change, and many sets of scenarios from different starting conditions can be generated using the same database.

To make this practical, we must define five processes:

- 1. A process to compute the yield curve given initial values of the state variables.
- 2. A process to generate a database of yield curves indexed by initial values of the state variables.
- 3. A process to retrieve a yield curve from the database based on the state variables, interpolating as needed when the values of the state variables do not exactly match any in the database.

- 4. A process to determine the values of the state variables that correspond to any specific yield curve. This is needed to initialize the state variables based on the starting observed yield curve.
- 5. A process to generate scenarios of yield curves using the database.

Each of these processes is described in a separate subsection below.

Completing the yield curve

The yield curve is to be completed based on this expectation:

$$P_T = E^Q \left[exp\left(-\int_0^T r_t dt \right) \right]$$

The expectation defines spot prices by term to maturity and can easily be converted to a curve of spot rates by term to maturity.

The expectation can be calculated as the simple average of many scenario-specific calculations using the path of r_t to calculate exp $\left(-\int_0^T r_t dt\right)$ for all values of T to get the spot prices. Since our scenarios are discrete rather than continuous, we use an approximation.

Define f = dt = the fraction of a year used as a time step. For monthly, $f = \left(\frac{1}{12}\right) = 0.083333$...

$$\exp\left(-\int_0^T r_t \, dt\right) \approx \prod_{t=1}^{T/f} \left(1 + \frac{r_{t-1} + r_t}{2}\right)^{-f}$$

On the right-hand side, t denotes the time step number and T/f is the total number of time steps. Note that we use the average of the short rate at the beginning and end of the period as the discount rate for the period.

To get a good approximation to the complete spot price curve, we use 50,000 scenarios for the short rate to estimate the expected value.

The yield curve based on these spot prices (using real-world parameters) does not include term premiums. One could easily add a fixed set of term premiums by maturity. In this model a more sophisticated approach is used. Term premiums, in theory, reflect the sensitivity of the spot price to changes in the state variables. That sensitivity can be measured if we have a database of these yield curves for different values of the initial state variables. The approach used here to generate the database of yield curves without term premiums is described in the next section. The section after that explains how the term premiums are estimated and added.

Generating the database

Specifying the contents

The database is simply a collection of yield curves indexed by the initial values of the state variables corresponding to each curve. One needs to specify which discrete values to include for each state variable, and then calculate and store yield curves for every possible combination of the three state variables.

To specify which discrete values to include for each state variable, specify the minimum value, the maximum value, and the size of the increment between each value and the next greater one. The total number of discrete values for each state variable is then

Number of discrete values = ((max – min) / increment) + 1

To get the total number of curves in the database, multiply together the number of discrete values for each state variable. For example, the following specification was used when working with the three-factor model described here:

State variable	Min	Max	Increment	No. values
Short rate	-0.05	0.25	0.005	61
Current mean	0.00	0.21	0.005	44
Volatility multiple	0	2	0.5	5

This resulted in a database of $61 \times 44 \times 5 = 13,420$ yield curves, conceptually in a 3-dimensional grid.

Judgment is needed in determining the increment and the number of different values needed for each state variable. In this case a small number of values was used for the volatility multiple because volatility tends to mean revert quickly and the initial volatility has a relatively small effect on the resulting yield curve compared to the effect of the other state variables.

Generating a database containing 13,240 yield curves can take some time. For this work, 50,000 scenarios were used in estimating each yield curve. On a workstation-class PC with 12 processor cores, database generation took about three hours. On a laptop with 4 processor cores, it took about eight hours. The author views this as practical because once the database is available, thousands of scenarios can be generated from any starting yield curve in a few minutes.

Calculating the yield curves

In a previous section on "Completing the yield curve" we explained how to simulate the yield curve given initial values for the state variables. But we stopped before adding the term premiums.

The process for adding the term premiums begins with generating the database of yield curves without term premiums. Term premiums are then added to each curve in the database by determining the sensitivity of spot prices to specific changes in the state variables. Calculating that sensitivity requires one to be able to retrieve yield curves from the database for different initial values of the state variables. The process for retrieving a curve for any set of initial state variables is described in the next subsection. Here we explain how to use that capability to calculate term premiums and update the database to include them.

As explained in reference [1] the term premium is theoretically equal to the following:

$$Term \ premium = \ \lambda \sigma_r \left(\frac{\partial P_T}{\partial r_0}\right)$$

 λ = market price of risk (a number generally in the range of 20% to 40% (0.20 to 0.40))

 P_T = spot price for maturity T

 r_0 = the initial short-term rate

 σ_r = volatility of the short-term rate

This formula applies to a one-factor short-rate model where the initial short rate is the only factor. In this case we have a three-factor model, and that makes it more complicated. To get a usable value, we conceptualize the formula to:

Term premium = λ x (change in spot price from a one standard deviation shock to state variables)

When there are three state variables each with their own volatility, calculating a one-standard-deviation random shock gets complicated. For our purposes a simple approximation is to calculate the effect on the state variables from a shock using a specific set of three standard normal deviates fed into the stochastic process. The three standard normal deviates correspond to $X_t^{(r)}$, $X_t^{(c)}$, $X_t^{(v)}$ in the formulas for the three stochastic processes in the model.

We begin with an uncorrelated set of random normal deviates (1,0,0) where the only nonzero shock is to the short rate. Since the model employs correlation between the short rate and the current mean reversion point, we approximate a correlated shock to the current mean equal to the correlation ρ thereby obtaining a set of shocks (1, ρ , 0) for ($X_t^{(r)}, X_t^{(c)}, X_t^{(v)}$). Those shocks are applied to the initial values of the state variables ($r_{t-1}, c_{t-1}, v_{t-1}$) to get a second set of values for the state variables (r_t, c_t, v_t) using the stochastic processes shown in Appendix 1.

With two sets of state variables derived in this way, we can use the database to retrieve the yield curve for each set of state variables. Then we calculate spot prices based on each yield curve and multiply the difference in spot prices by λ to get the term premiums.

The yield curve for a set of state variables is a set of values $r_{t,T}$ for selected maturities (T). The spot prices are $P_{t,T} = (1 + r_{t,T})^{-T}$. The two yield curves give two sets of spot prices, and the difference between the two spot prices for maturity T can be called $\Delta P_{t,T}$. If the two sets of state variables differ by one monthly standard deviation, then $\Delta P_{t,T}$ represents the standard deviation of the price over a period of one month. We approximate the standard deviation over one year by multiplying the standard deviation over one month by the square root of 12. Then we multiply by λ (market price of risk) to get the term premium $\varpi_{t,T}$.

$$\varpi_{t,T} = \Delta P_{t,T} \sqrt{12}\lambda$$

The term premiums are then added to the yield curve retrieved using the initial values for the state variables. Since that yield curve did not contain term premiums (it was retrieved from a database of curves without term premiums) the adjusted curve with term premiums added is the curve we really want in the database. This process is carried out for each curve in the database without term premiums and a new database of curves with term premiums is created and saved. The database of curves without term premiums is discarded.

Retrieving a yield curve

We need to be able to retrieve a yield curve when given specific values for the three state variables. This is done by retrieving some curves from the database and interpolating between them.

The database has curves only for specific values of each state variable. Since those values are at equal intervals, a first step is to determine the interval in which the specified value of the state variable resides. Determine the values at the top and the bottom of that interval. Yield curves are available for those values. Do this for all three state variables, obtaining the top and bottom value for each. Then construct eight sets of state variables, each set containing either the top or bottom value for each state variable, in every combination.

	Var 1	Var 2	Var 3
1	Тор	Тор	Тор
2	Тор	Тор	Bottom
3	Тор	Bottom	Тор
4	Тор	Bottom	Bottom
5	Bottom	Тор	Тор
6	Bottom	Тор	Bottom
7	Bottom	Bottom	Тор
8	Bottom	Bottom	Bottom

Then retrieve a yield curve from the database for each of the eight sets of state variables. The desired yield curve is obtained through three-dimensional linear interpolation between these eight, based on the exact specified values for the state variables.

To perform three-dimensional linear interpolation, we first determine where each state variable is within its interval. For each state variable, calculate the fraction $f = \frac{actual \ value - bottom \ of \ interval}{top \ of \ interval - bottom \ of \ interval}$ to get three values f1, f2, and f3. These fractions are used to calculate a weight for each of the eight yield curves. The weights sum to 1.0 and the weighted sum is the interpolated value. Here are the weights:

	Weight
1	f1 x f2 x f3
2	f1 x f2 x (1-f3)
3	f1 x (1-f2) x f3
4	f1 x (1-f2) x (1-f3)
5	(1-f1) x f2 x f3
6	(1-f1) x f2 x (1-f3)
7	(1-f1) x (1-f2) x f3
8	(1-f1) x (1-f2) x (1-f3)

Initializing the state variables

State variables are initialized to fit an initial yield curve by finding the set of state variables which, when used to retrieve a curve from the database, provides the best fit curve.

Goodness-of-fit can be defined in many ways. We choose to minimize the sum of squared differences between the desired initial curve and the curve from the database. We use the 10 points on the yield curve provided in the Federal Reserve Board's H 15 release. Take 10 differences, square them, and sum the squares to get the goodness-of-fit statistic.

This amounts to a nonlinear fitting exercise, and such exercises are prone to find local minimums. To prevent that, a two-step process is used. First there is a global search through every curve in the database to find the state variables corresponding to the curve with the best fit. The desired state variables must be close to those, so a gradient search is then performed in the neighborhood around those state variables to arrive at the best fit. The gradient search uses a much smaller interval between state variable values than the increment between state variable values that appear in the database, and therefore represents a search over interpolated curves near the best fitting curve in the database.

The gradient search uses an increment between state variable values that is 1/20th the size of the interval between values in the database. The search proceeds by tweaking the state variables one at a time by that small increment, as follows:

- 1. For state variable 1
 - a. Add the small increment to the value.
 - i. Calculate a new yield curve using the adjusted set of state variables.
 - ii. Calculate the goodness of fit of the new yield curve.
 - iii. If the goodness of fit is improved, then set state variable 1 to the new value, record that it was changed, and go to step 2.
 - b. If the goodness of fit was not improved in step 1a, subtract a small increment from the value.
 - i. Calculate a new yield curve using the adjusted set of state variables.
 - ii. Calculate the goodness of fit of the new yield curve.
 - iii. If the goodness of fit is improved, then set state variable 1 to the new value, record that it was changed, and go to step 2.
- 2. Same as step 1 but adjust the value of state variable 2 by its small increment.
- 3. Same as step 2 but adjust the value of state variable 3 by its small increment.
- 4. If any of the state variables were changed in steps 1, 2, or 3 then go back to step 1 and continue the gradient search. Otherwise, the search is completed, and the best fitting set of state variables has been found.

Generating scenarios using the database

The process of generating scenarios using the database makes use of the processes described above. The steps are:

- Initialize the state variables based on the starting yield curve. Note that this includes initialization of the current mean reversion point, so the current mean reversion point cannot be specified ahead of time as a constant.
- 2. Use the stochastic process to create scenario paths for the set of three state variables.
- 3. At each time step, retrieve a yield curve from the database based on the values of the state variables for that time step.

During scenario generation curves are retrieved from the database rather than being simulated numerically. This makes the scenario generation process reasonably fast.

Appendix 1: Summary of stochastic processes in this model

In this appendix we define the stochastic process used for each of the state variables r_t , c_t , v_t used in this model. Several kinds of parameters such as mean reversion point and volatility are used in each of these processes, so we use superscripts in our notation to indicate the state variable to which each parameter applies, as in $\sigma^{(r)}$, $\sigma^{(c)}$, $\sigma^{(v)}$ for the volatility parameters.

Volatility process

The state variable v_t follows this stochastic process:

$$v_t = v(r_{t-1})e^{x_{t-1}-1}$$

 $v(r_{t-1})$ is the tabulated volatility for short rate r_{t-1} interpolated between tabulated values.

 x_t is a volatility multiplier following the stochastic process:

$$x_{t} = x_{t-1} + \gamma^{(v)} (\mu^{(v)} - x_{t-1}) + \sigma^{(v)} X_{t}^{(v)}$$

Parameters:

 $\mu^{(v)}$ = Mean reversion point for the volatility multiplier (normally 1.0 or near 1.0)

 $\sigma^{(v)} =$ Volatility of the multiplier

 $\gamma^{(v)}$ = Mean reversion strength for the multiplier

Random element at each time step:

 $X_t^{(v)}$ = Random shock (standard normal distribution, mean 0, variance 1)

Mean reversion point process

The state variable c_t follows this stochastic process:

$$c_{t} = \max\{f, c_{t-1} + \gamma^{(c)} (\theta^{(c)} - c_{t-1}) + \sigma^{(c)} X_{t}^{(c)}\}$$

 $\theta^{(c)} =$ a fixed long-term mean around which the moving mean fluctuates.

 $\gamma^{(c)}$ = the rate at which the moving mean is attracted to (reverts to) the fixed long-term mean.

 $\sigma^{(c)} =$ volatility from a table of values indexed by c_{t-1} and interpolated as needed.

 $f = \text{floor} = \max(\psi, 2r_{t-1} - 0.34)$ where ψ is the floor in effect when the short rate is low.

 $X_t^{(c)} =$ Random shock (standard normal distribution, mean 0, variance 1)

Short rate process

The short rate r_t follows this stochastic process, which depends on values of all three state variables from the prior time step. The process is defined using a drift function and a random shock function:

$$r_t = r_{t-1} + d(r_{t-1}, c_{t-1}) + s(r_{t-1}, v_{t-1})$$

The drift function is: $d(r_{t-1}, c_{t-1}) = \gamma^{(r)} \eta(c_{t-1} - r_{t-1})$

 $\gamma^{(r)}$ = the typical rate at which the short rate is attracted to (reverts to) the current mean.

 η = a conditional multiplier whose value depends on the relationship between r_{t-1} , c_{t-1} , $\theta^{(c)}$ as follows:

We define two parameters:

$$Z_h =$$
 high stretch, value ≥ 1.0

$$Z_l =$$
low stretch, value ≥ 1.0

The multiplier η takes one of the following values:

If
$$c_{t-1} < \theta^{(c)}$$
 then
If $r_{t-1} > c_{t-1}$ then $\eta = Z_l$ else $\eta = 1/Z_l$
If $c_{t-1} > \theta^{(c)}$ then
If $r_{t-1} < c_{t-1}$ then $\eta = Z_h$ else $\eta = 1/Z_h$

The random shock function is: $s(r_{t-1}, v_{t-1}) = v_{t-1}X_t^{(r)}$

 v_{t-1} = The state variable for the beginning of month stochastic volatility (process described above) $X_t^{(r)}$ = Random shock (standard normal distribution, mean 0, variance 1)

Updating the state variables for each time step

To update the state variables for each time step we begin by generating three independent random shocks $X_t^{(r)}, X_t^{(c)}, X_t^{(v)}$, each from a Gaussian Normal distribution with mean 0 and variance 1. Then, since we need $X_t^{(c)}$ to be correlated with $X_t^{(r)}$ we modify the value of $X_t^{(c)}$:

$$New \, X_t^{(c)} = \rho X_t^{(r)} + \sqrt{1-\rho^2} (unmodified \, X_t^{(c)})$$

The revised set of values $X_t^{(r)}$, $X_t^{(c)}$, $X_t^{(v)}$ can then be used to update each of the state variables using the formulas for the stochastic process of each state variable.

Appendix 2: Suggested calibration and illustrative scenario paths

Adjusting for the discrete time step

When a model is expressed in discrete form rather than continuous form, one must specify the size of the time step for which parameter values are specified. Parameter values for volatility and for mean reversion rate depend on the size of the time step in the following way:

Let n = number of time steps per year. For monthly time steps, n = 12.

 $Discrete \ volatility = \frac{Annual \ volatility}{\sqrt{n}}$

Discrete mean reversion rate = $1 - e^{-(Annual mean version rate)/n}$

The annual mean reversion rate in this formula is defined with reference to an exponential decay process in a continuous model. If the annual mean reversion rate is γ , then in a continuous model the fraction of the distance between the current value and the mean covered by mean reversion in one year is $1 - e^{-\gamma}$. If there are *n* time steps per year, then the fraction of that distance covered in each time step is $1 - e^{-\gamma/n}$.

Parameter values shown in this appendix are for a discrete one-month time step, where n = 12.

Volatility process

The state variable v_t follows this stochastic process:

$$v_t = v(r_{t-1})e^{x_{t-1}-1}$$

 $v(r_{t-1})$ is the tabulated volatility for short rate r_{t-1} interpolated between tabulated values.

 x_t is a volatility multiplier following the stochastic process:

$$x_{t} = x_{t-1} + \gamma^{(v)} (\mu^{(v)} - x_{t-1}) + \sigma^{(v)} X_{t}^{(v)}$$

Parameter values:

 $\mu^{(v)} = 1.0$

 $\sigma^{(v)} = 0.072$ (annualized value is 0.25)

 $\gamma^{(\nu)} = 0.072$ (annualized value is 0.9)

v(r) is linearly interpolated between these values:

r	v(r)
-0.2% or below	0.0
2%	0.48%
6%	0.48%
14% or above	1.50%

Note: The tabulated value of 0.48% at r = 2% is near the top of the range of observed volatility when interest rates have been around 2%, as shown by the red line in the chart below. The suggested value was developed based on fitting the short rate process to many different measures of historical behavior.

Moving that value up from the center of the observed range improved the fit of the model in other areas, notably the distribution of yield curve slopes when interest rates are low.



Mean reversion point process

The state variable c_t follows this stochastic process:

$$c_{t} = \max\{f, c_{t-1} + \gamma^{(c)} (\theta^{(c)} - c_{t-1}) + \sigma^{(c)} X_{t}^{(c)}\}$$

Parameter values:

$$f = \max(\psi, 2r_{t-1} - 0.34)$$

 $\psi = -0.002$

 $\theta^{(c)} = 0.04$

 $\gamma^{(c)} = 0.00416$ (annualized value is 0.05)

$$ho = 0.2$$
 (used to obtain $X_t^{(c)}$)

 $\sigma^{(c)} =$ volatility from a table of values indexed by c_{t-1} and linearly interpolated as needed.

С	$\sigma^{(c)}$
002 or below	0.0
0.02	0.003
0.06	0.004
0.14 or above	0.002

The floor and the tabulated volatility were introduced to make the emergent behavior of the short rate process more like that historically experienced. To understand why, consider the distribution of the mean reversion point 50 years in the future.

If there were no floor and the volatility were constant at 0.003, the distribution would look like this:



The floor was introduced because it seemed unreasonable for the mean reversion point to go far below zero. The distribution with a floor but constant volatility would look like this:



The volatility was modified from a constant to a value tabulated by the current mean, being higher when the current mean is near its average and lower when the current mean is far from average. This stretched the middle of the distribution shown above without stretching the tails. The result looks like this:



The moving mean reversion point is not an observable variable, so one cannot compare this distribution with observed historical values. It can only be evaluated indirectly through its effect on the emergent distribution of the short rate process that depends upon it, and through its effect on the ability of the model to fit historical yield curve shapes. Nevertheless, one can speculate that the reason why this distribution helps emulate historical behavior is because the Fed tends to lean towards employment stimulus (lower interest rates) more often than towards inflation control (higher interest rates).

Short rate process

The short rate r_t follows this stochastic process, which depends on values of all three state variables from the prior time step. The process is defined using a drift function and a random shock function:

$$r_t = r_{t-1} + d(r_{t-1}, c_{t-1}) + s(r_{t-1}, v_{t-1})$$

The drift function is: $d(r_{t-1}, c_{t-1}) = \gamma^{(r)} \eta(c_{t-1} - r_{t-1})$

Parameter values:

 $\gamma^{(r)} = 0.0206$ (annualized value is 0.25)

 η = a conditional multiplier whose value depends on the relationship between r_{t-1} , c_{t-1} , $\theta^{(c)}$ as follows:

We define two parameters:

$$Z_h = 0.75$$
 (high stretch)

 $Z_l = 3.0$ (low stretch)

The multiplier η takes one of the following values:

If $c_{t-1} < \theta^{(c)}$ then If $r_{t-1} > c_{t-1}$ then $\eta = Z_l$ else $\eta = 1/Z_l$ If $c_{t-1} > \theta^{(c)}$ then If $r_{t-1} < c_{t-1}$ then $\eta = Z_h$ else $\eta = 1/Z_h$

The random shock function is: $s(r_{t-1}, v_{t-1}) = v_{t-1}X_t^{(r)}$

Parameter values for the volatility process v_t used in the random shock function were given earlier in this Appendix.

Completing the yield curve:

The market price of risk used to add term premiums when completing the yield curve is:

 $\lambda = 0.20$

Some interest rate scenario paths based on this calibration are shown below. The scenarios cover 50 years, and the charts show monthly values for the 1-year rate and the 20-year rate from the yield curve.

This scenario illustrates some of the stylized facts. The short-term rate is more volatile than the long-term rate. When the short-term rate spikes upwards, the slope of the yield curve (20yr rate minus 1yr rate) tends to decline and the curve may become inverted (1yr rate > 20yr rate).



The scenario below illustrates a "low for long" scenario with the short-term rate stuck near zero for an extended period, during which the long-term rate continues to fluctuate.



The chart below illustrates the distribution of yield curve slopes. Each point represents a yield curve with the 1yr rate plotted against the 20yr rate. The diagonal line shows where the two rates would be equal. Observe that:

- When the 1yr rate is under 2% the slope of the yield curve is virtually always positive (20yr rate greater than 1yr rate)
- While negative slopes do occur when interest rates are over 2%, there is a wide distribution of slopes even when interest rates are very high.



The shape of the "blob" on this plot is very much like the "blob" one gets when plotting historical yield curves in the same way, although the number of historical data points is smaller, especially at high interest rates.

References

[1] <u>Strommen, Stephen J., "Understanding the Connection between Real-World and Risk-Neutral</u> <u>Scenario Generators", SOA Research Institute, 2022</u>